

ABSTRACT

A computer system having a fault-tolerance framework in an extendable computer architecture. The computer system is formed of clusters of nodes where each node includes computer hardware and operating system software for executing jobs that implement the services provided by the computer system. Jobs are distributed across the nodes under control of a hierarchical resource management unit. The resource management unit includes hierarchical monitors that monitor and control the allocation of resources. In the resource management unit, a first monitor, at a first level, monitors and allocates elements below the first level. A second monitor, at a second level, monitors and allocates elements at the first level. The framework is extendable from the hierarchy of the first and second levels to higher levels where monitors at higher levels each monitor lower level elements in a hierarchical tree. If a failure occurs down the hierarchy, a higher level monitor restarts an element at a lower level. If a failure occurs up the hierarchy, a lower level monitor restarts an element at a higher level. Each of the monitors includes termination code that causes an element to terminate if duplicate elements have been restarted for the same job. The termination code in one embodiment includes suicide code whereby an element will self-destruct when the element detects that it is an unnecessary duplicate element.